

110826

31798

p. 6

# DESIGNING AN AUTONOMOUS ENVIRONMENT FOR MISSION CRITICAL OPERATION OF THE EUVE SATELLITE

Annadiana Abedini, Roger F. Malina  
Center for EUV Astrophysics  
University of California  
2150 Kittredge St.  
Berkeley, CA 94720-5030

**Abstract**—Since the launch of NASA's *Extreme Ultraviolet Explorer (EUVE)* satellite in 1992, there have only been a handful of occurrences that have warranted manual intervention in the *EUVE* Science Operations Center (ESOC). So, in an effort to reduce costs, the current environment is being redesigned to utilize a combination of off-the-shelf packages and recently developed artificial intelligence (AI) software to automate the monitoring of the science payload and ground systems. The successful implementation of systemic automation would allow the ESOC to evolve from a seven day/week, three shift operation, to a seven day/week one shift operation.

First, it was necessary to identify all areas considered *mission critical*. These were defined as:

- The telemetry stream must be monitored autonomously and anomalies identified.
- Duty personnel must be automatically paged and informed of the occurrence of an anomaly.
- The "basic" state of the ground system must be assessed. Monitors should check that the systems and processes needed to continue in a "healthy" operational mode are working at all times. Network loads should be monitored to ensure that they stay within established limits.
- Connectivity to Goddard Space Flight Center (GSFC) systems should be monitored as well: not just for connectivity of the network itself but also for the ability to transfer files.
- All necessary peripheral devices should be monitored. This would include the disks, routers, tape drives, printers, tape carousel, and power supplies.
- System daemons such as the archival daemon, the Sybase server, the payload monitoring software, and any other necessary processes should be monitored to ensure that they are operational.
- The monitoring system needs to be redundant so that the failure of a single machine will not paralyze the monitors.
- Notification should be done by means of looking through a table of the pager numbers for current "on call" personnel. The software should be capable of dialing out to notify, sending email, and producing error logs.
- The system should have knowledge of when real-time passes and tape recorder dumps will occur and should know that these passes and data transmissions are successful.

Once the design criteria were established, the design team split into two groups: one that addressed the tracking, commanding, and health and safety of the science payload; and another group that addressed the ground systems and communications aspects of the overall system.

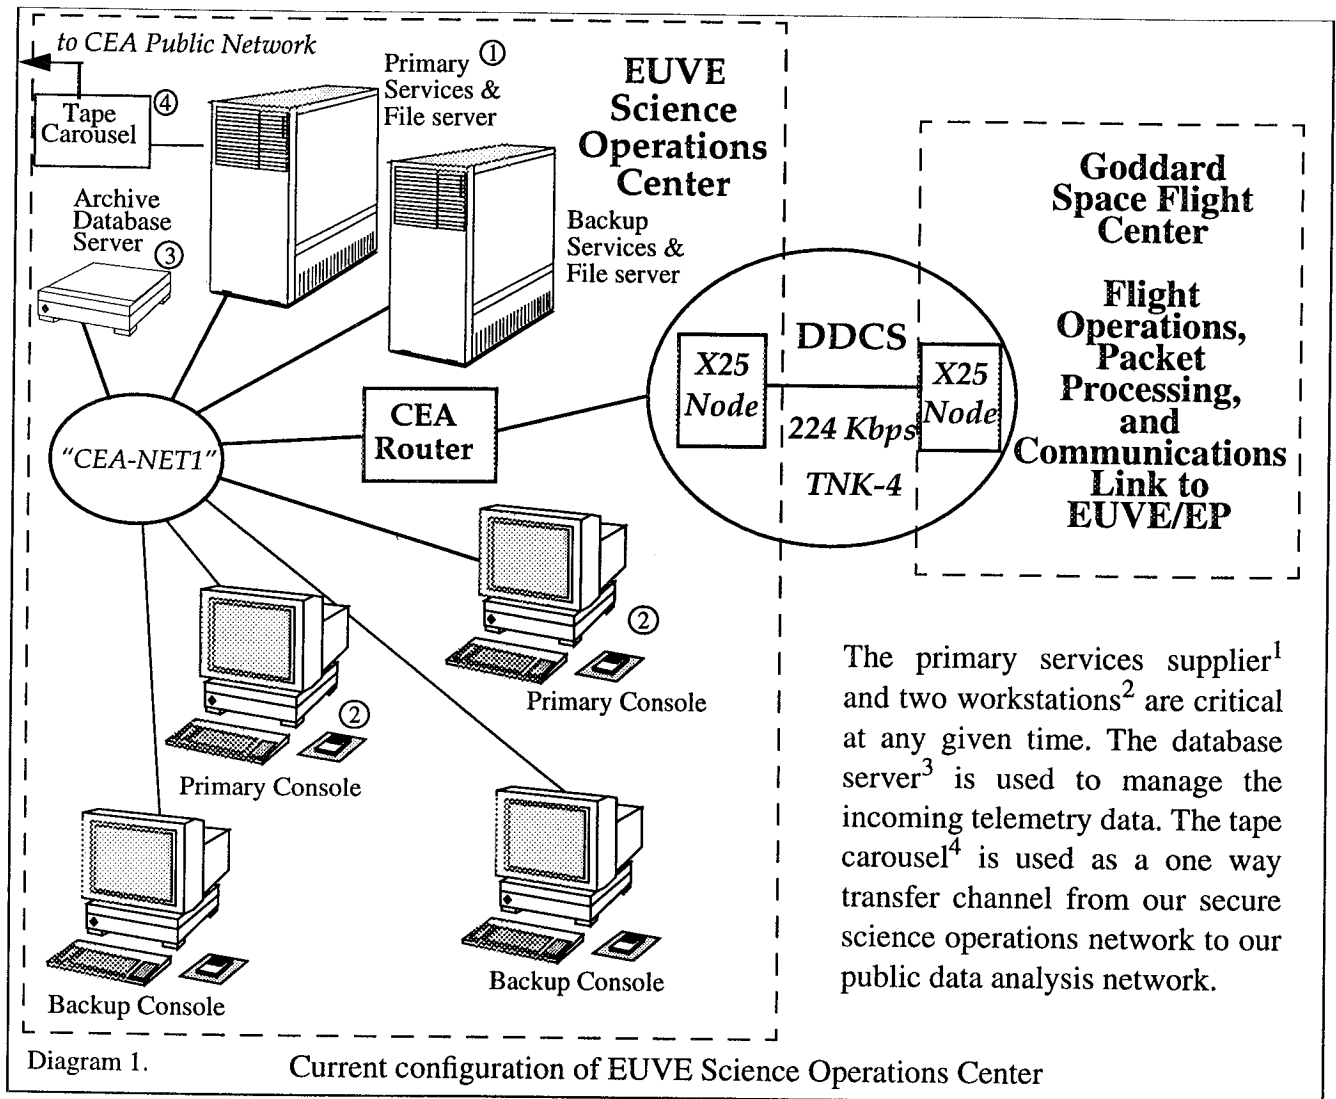
## INTRODUCTION

In June, 1992, NASA launched the *Extreme Ultraviolet Explorer (EUVE)* satellite (Bowyer & Malina, 1991). The science payload for *EUVE* was designed and built at the University of California, Berkeley. The operations center (ESOC) for the science payload is located at the Center for EUV Astrophysics (CEA) at UC Berkeley.

The current method used for monitoring the *EUVE*'s science payload is a program called "soctools," which was developed at CEA. This tool displays numerical tables that change color and, in some cases,

give audible output when a monitored value goes out of or back into limit. Soctools is dependent upon the presence of a human who watches the printed displays as the values change. This situation requires 24 hours per day, 7 days per week staffing of the ESOC.

Though many aspects of the operations have been automated to some extent,<sup>1</sup> the monitoring of the EUVE/Explorer Platform (EP) science payload, as well as the monitoring and reconfiguration of the ground systems of CEA's secure science operations network,<sup>2</sup> is done manually.



The primary services supplier<sup>1</sup> and two workstations<sup>2</sup> are critical at any given time. The database server<sup>3</sup> is used to manage the incoming telemetry data. The tape carousel<sup>4</sup> is used as a one way transfer channel from our secure science operations network to our public data analysis network.

At the time that the ESOC was originally designed, the systems and software, which could allow a robust, fault-tolerant computing environment, were prohibitively expensive. With the desire to acquire a 99% delivery of data to CEA, the current configuration of the ESOC contains redundant computing and network hardware to reach this goal at a more reasonable cost. In the case of a failure of any of the integral parts of this system, it is possible for the payload controller or a hardware support person to

1. The telemetry reception and storage as well as the transfer of the telemetry from the ESOC network to a mass-storage optical jukebox on the science data analysis network are all done by means of automated software programs.
2. CEA's public science data analysis systems and network are physically separate from the science operations network.

remove a failed device from operation and quickly introduce a similar device into operation in its place. In the above diagram, the *services supplier* is responsible for providing the processes that monitor the reception and storage of incoming telemetry data. The services supplier also provides the localized storage space used by the system utilities, login directory space needed by the payload controllers and other ESOC supporting staff, and temporary storage for the incoming telemetry.

The *backup services supplier* stands ready to be reconfigured if a problem arises with the primary system. The backup services supplier is also available for development and data analysis when it is not needed in the role of primary services supplier.

Similarly, the configuration of the database server has been planned such that one of the *backup console* workstations can be easily reconfigured in case of a controlling CPU failure. However, the tape carousel, which has proven to be one of our weakest links, has no *hot spare*. For this reason, we have set aside enough disk space to store several days of telemetry. In a situation where the carousel is down for a longer period of time than the disk space could accommodate, there is even a *sneaker-net*<sup>1</sup> plan that outlines a procedure for moving the telemetry from the "CEA-NET1" to our public network by hand, if necessary.

As noted above, after the statement of criteria had been defined, the initial design team was split into two parts, one to focus on the ground systems and communications, the other to focus on the monitoring of the science payload. The latter of these groups later divided the design efforts into health and safety monitoring of the payload, and the commanding of the of the payload. The team has postponed addressing, in depth, a design plan that would focus on the automation of commanding.<sup>2</sup>

## THE SCIENCE PAYLOAD

The team evaluated several commercial and NASA funded packages intended to monitor satellite telemetry. The software package RTworks® was chosen. This commercial package contains generalized tools that can be customized to fit the specific needs of a project. It contains tools for building user displays as well as the capability to be "taught" about processes and to initiate other processes when an anomaly has been detected.

Initially, the team selected a set of six critical engineering monitors for autonomous monitoring using RTworks. The payload controllers captured their procedural knowledge into flow charts, which were then transformed into data-flow diagrams,<sup>3</sup> as the controllers worked with a small team of programmers.

The lengthy process of creating and reworking these diagrams requires many iterations. However, the resulting diagrams not only give an accurate representation of the detection process but provide the programmers with an accurate starting point.

After both the controller and the programmer were satisfied with the visual representation of the anomaly detection process for each of the monitors in question, the result was programmed into RTworks' inference engine, creating EUVE specific extensions we call Eworks.

As the telemetry is received, Eworks actively monitors the telemetry data stream and if an anomaly is detected, a process is initiated which will notify someone of the occurrence. Currently this is done by means of initiating a program which pages the on-call controller.

Eventually, we would like to add diagnostic capabilities that would then allow autonomous response to

- 
1. Sneaker-net is an industry buzz word that describes the process of hand carrying information from one location to another, or from one system to another, usually on a tape or floppy disk.
  2. The team hopes to resume efforts in the area of automated commanding once confidence has been gained in the use of the monitoring software.
  3. Flow charts may be adequate for the documentation of step by step procedures, but are not always the best means for representing a detailed picture of a complex, interactive system.

some of the anomalies. Since the software contains the ability to initiate other programs when an anomaly is detected, having the software invoke a diagnostic process that would then initiate a corrective action sequence, would be the logical next step.

## **THE SCIENCE OPERATIONS CENTER GROUND SYSTEMS**

### **Monitoring**

The team also investigated several software packages that could be used to monitor the availability and capacity of disks, system and network services, and critical processes. The team selected the commercial software package Sun NetManager® (SNM) for use for this task.

Like RTworks, in addition to its monitoring capabilities, SNM possesses the ability to initiate a corrective action sequence when an anomaly is detected. This means that SNM can be configured to either notify someone of the anomaly or take corrective action as appropriate.

The software can be configured so that corrective action is taken when disk usage of critical areas exceeds a specified limit, or when critical system or network services become unavailable. In example, if the primary services supplier becomes unavailable, the SNM software can start the lost services on the backup services supplier without human intervention. It can also initiate a process that will page the on-call hardware person and notify them of the loss of the primary server.

Similarly, if a critical disk area is filled above a prescribed limit, SNM can initiate a program that will clean up the area and remove files that are no longer needed.

### **Systems configuration**

Although the automated software can simplify the overall monitoring process, the interdependence of the systems and peripherals are still a point of failure that would require human intervention in many cases. For this reason, suggestions have been made that would reduce these dependencies.

The current configuration utilizes multi-processor computer servers that share the tasks of providing disk storage to the systems on the network as well as processing power. The recent introduction of a networked version<sup>1</sup> of the redundant array of independent disks<sup>2</sup> (RAID) disk technology allows us to resolve the problem of having to duplicate disk storage areas and user accounts on both of the services suppliers.

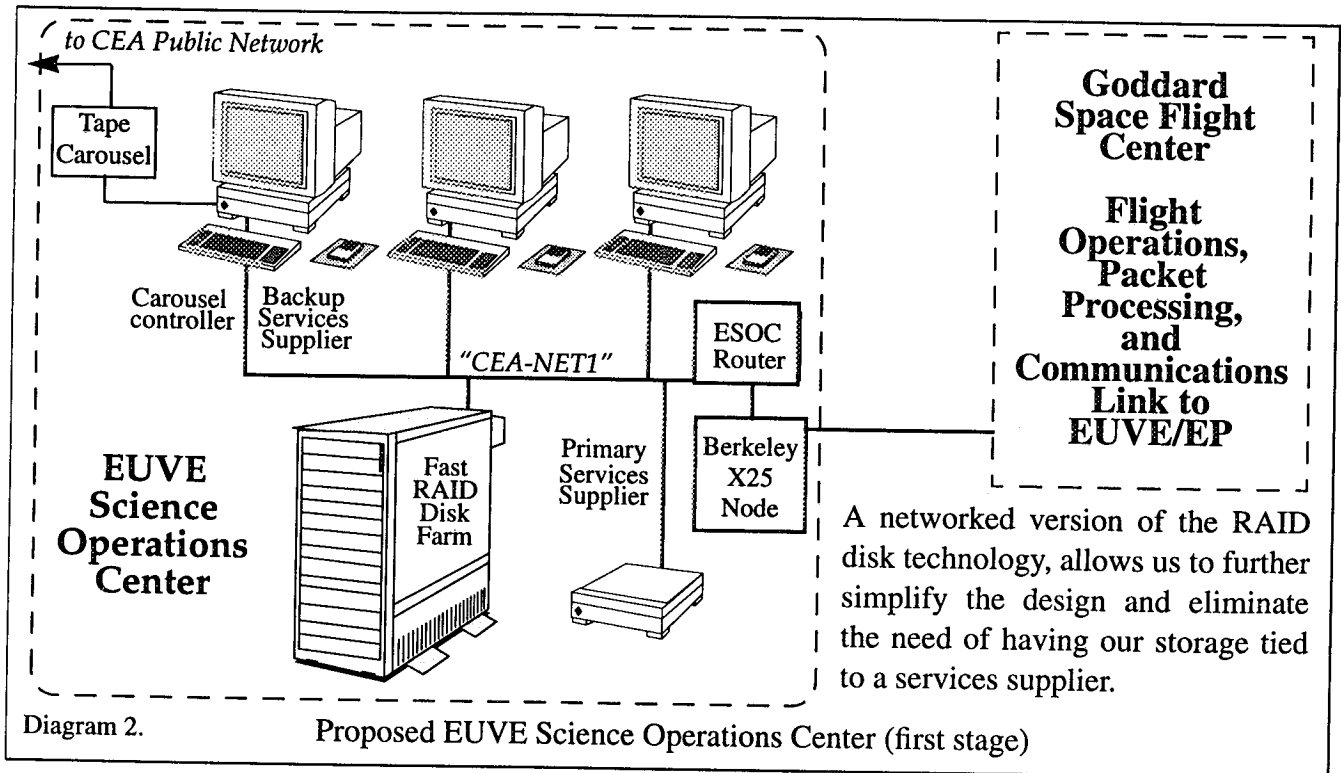
The RAID disk array provides highly reliable, fast disk storage to all of the systems on the network. With the utilization of disk stripping, a *warm spare* disk and a backup power supply, the system provides a hands-off, fault-tolerant storage solution. This unit also allows us to resolve many of the issues that set the requirement for human intervention in the transition from one disk and services supplier to another.

To date, our experience has shown that when there is a problem with the file server, which is currently supplying services to the ESOC, most of the systems in ESOC lock up and often require rebooting. This happens when the disk space that was being supplied by the services supplier becomes unavailable.<sup>3</sup> Since this RAID disk is attached to the network and not any single system, the RAID disk eliminates the need to reboot systems. If a system supplying services fails, the disk area is still available to all of the other systems on the network. Thus, there is no network lockup and no interruption of services.

The elimination of the network deadlock allows the monitoring SNM software to initiate a program that

- 
1. The FAServer 1400 from NAC is the first Unix RAID disk array that is not tied to another CPU.
  2. The description of RAID disk technology is beyond the scope of this paper. Please contact your hardware vendor for more specific information regarding this technology.
  3. If the disk was being actively accessed and the server does not come back on line, the workstation is locked into a disk-wait state.

will autonomously start any lost services on an alternate system and to dial out to notify hardware or software support personnel of the occurrence.



Having freed ourselves of the need for file server class systems, we can consider moving our processing needs to some of the fairly inexpensive, high powered, multi-processor workstation class computers. A computing system such as the Sun SparcStation 10 can supply an adequate level of computing power for the tasks of providing the daemons that oversee the reception and storage of telemetry, and monitoring the ground systems and science payload.

With the purchase of additional network cards, the NAC FAServer, RAID disk array, can allow up to four simultaneous network connections in order to supply the disk space to those networks. It, however, does not route information *between* those networks. The FAServer simply allows the users on all networks to be able to view the data stored on the disk array. The FAServer also allows for restricting access from a given network if desired. This means that the RAID disk array can be located in the access restricted computer room with direct login allowed only from its console. All other access to the unit would only be to the drives directly and the information stored there. This access could be restricted to read-only if desired.

This in mind, we could further simplify the configuration and allow the elimination of the tape carousel if we utilize the network RAID box feature that allows it to span multiple networks. In this alternate configuration the telemetry data, which is written to the disks as it comes in from the satellite, can be made available to the science data network as read-only information, and there would be no direct access to the secure network, nor opportunity to alter the data being provided by that network.

## CONCLUSION

As confidence is gained in both the hardware and software being introduced into the ESOC, we will relax the staffing requirements, which are currently needed to ensure a smooth running environment.

Additionally, CEA is currently involved in the testbedding of diagnostic software packages from Jet Propulsion Labs and NASA Ames Research Center, as stated earlier, that will facilitate the autonomous resolution of predictable anomalies. Once the anomaly diagnostic systems are mature, we hope to start utilizing these techniques in the diagnosis of detected anomalies as well as add autonomous resolution of diagnosed problems.

## ACKNOWLEDGMENTS

We thank M. Montemerlo, P. Friedland, D. Korsmeyer, and D. Atkinson for their collaboration on the work of new technology infusion into the EUVE Project. We would also like to thank the members of the CEA staff who are involved in the design and implementation of these operations improvements.

The work described here is funded by NASA contract NAS5-29298 and NASA Ames contract NCC2-838.

## REFERENCES

- Bowyer, S., & Malina, R. F. (1991). The EUVE Mission. *Extreme Ultraviolet Astronomy*, 397–408.
- RTworks, Talarian Corporation, 444 Castro St., Suite 140, Mt. View, CA 94041, (415) 965–8050.
- Sun NetManager, Sun Connect, 2550 Garcia Ave., Mt. View, CA 94043, (800) 241–CONX.
- FAServer 1400, Network Appliance, Corp., 295 N. Bernardo Ave., Mt. View, CA 94043, (415) 428–5100.
- SPARCsystem 10, Sun Microsystems Computer Corporation, 2550 Garcia Ave., Mt. View, CA 94043, (415) 960–1300.